IA 32 Processor Architecture Trends and Research

Justin Rattner

Intel Senior Fellow and Director
Microprocessor Research Labs

and

Joe Schutz

Vice President and Director
Logic Technology Development

Intel Corporation

Feb 19, 2003

Agenda

- Microprocessor Trends and Challenges
- Prescott: A State-of-the-Art Microprocessor on 90nm
- Microprocessor Research Directions
- Summary

"Iron Law" of Microprocessor Performance

$$\text{Processor Performance} = \frac{\text{Instructions}}{\text{Time}}$$

$$\frac{\text{Instructions}}{\text{Second}} = \frac{\text{Instructions}}{\text{Clock}} \times \frac{\text{Clocks}}{\text{Second}}$$

(inst. rate) (IPC) (frequency)

IPC = f(microarchitecture, process)

Frequency = f(process, circuits)

Frequency & Performance Boost

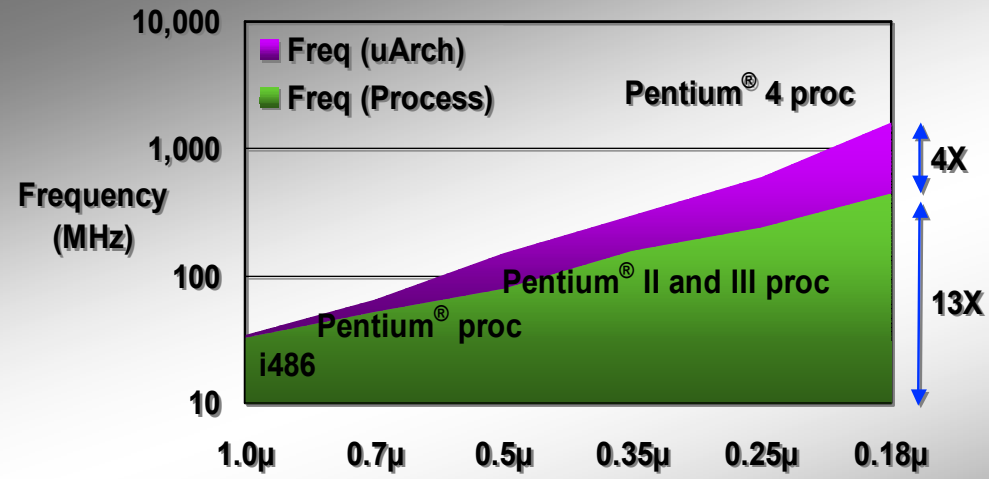

Frequency Increased 50X

- 13X due to process technology
- Additional 4X due to circuits

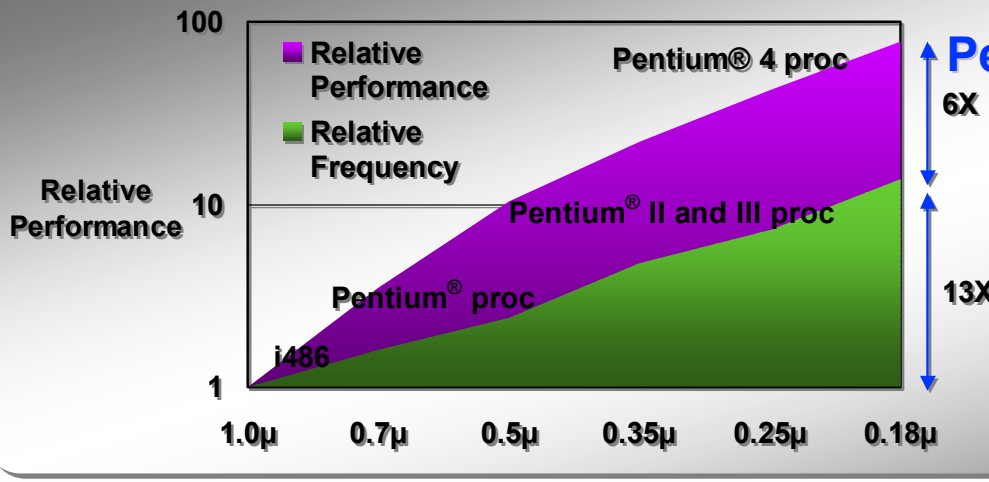

Performance Increased >75X

- 13X due to process technology
- Additional >6X due to microarchitecture

*Note: Performance measured using SpecINT and SpecFP

50X in frequency and 75X in performance

Source: Intel Corporation

Technology Scaling Theory

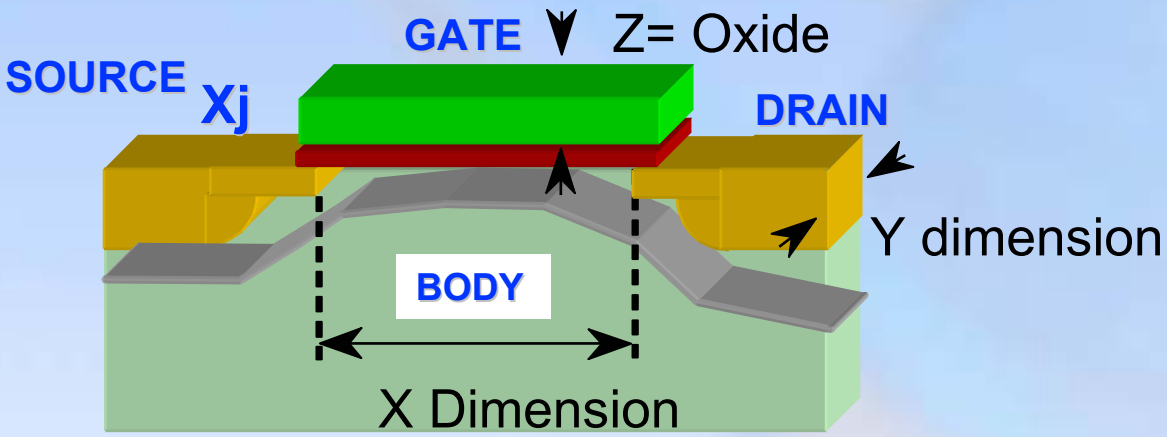

X & Y Dimensions scale down by 30%	**Doubles transistor density**
Z-Oxide thickness scales down	**Faster transistor, higher performance**
Vcc & Vt scaling	**Lower active power**

Technology scaling is a great thing

Technology Scaling Trends

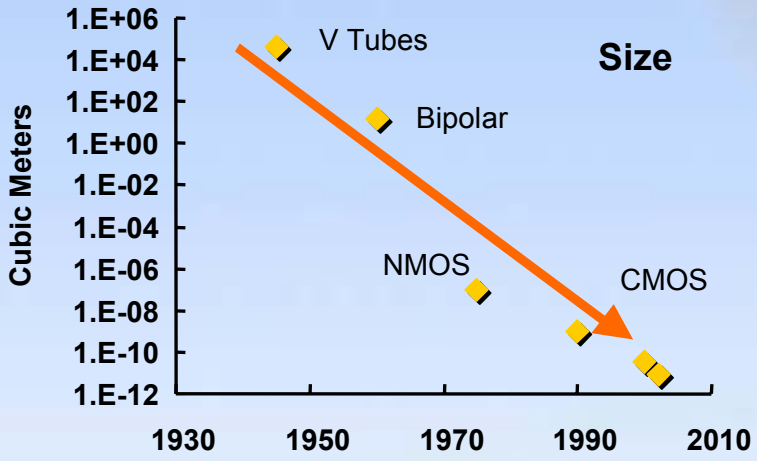

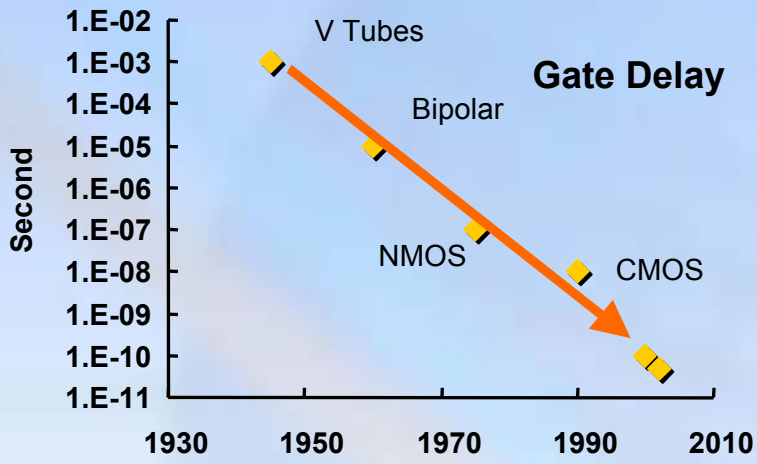

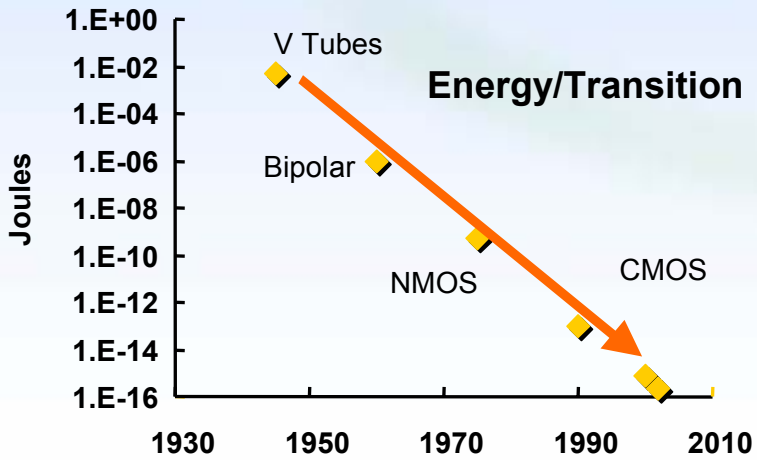

Technology Transitions

V Tubes ⇒ Bipolar
Bipolar ⇒ NMOS
NMOS ⇒ CMOS
CMOS ⇒ ?

intel

Scaling Will Continue

Supply Voltage Scaling Trends

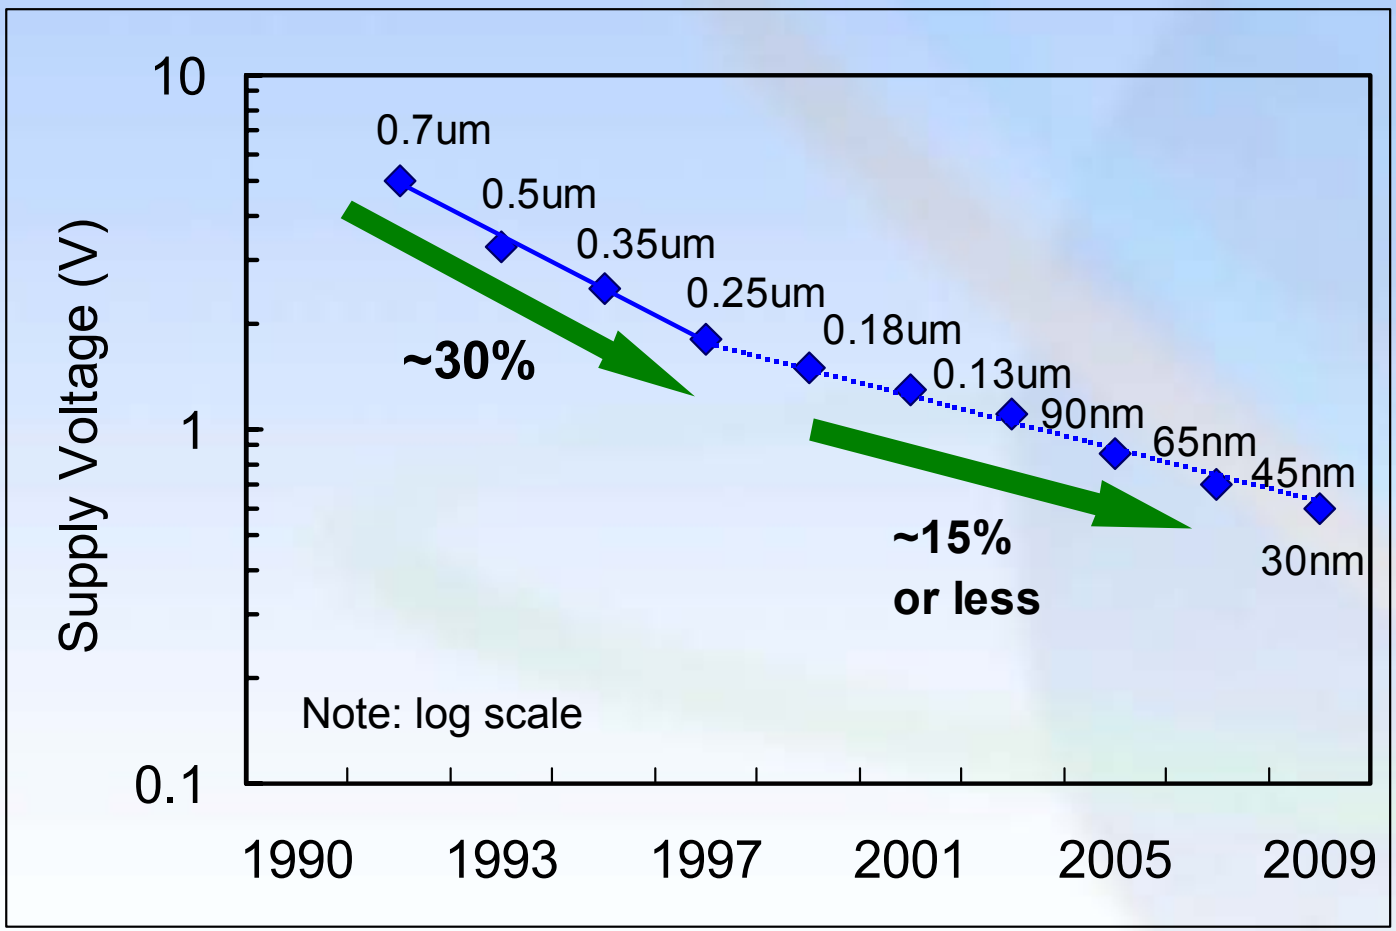

Voltage scaling has slowed

Growing Interconnect RC Delays

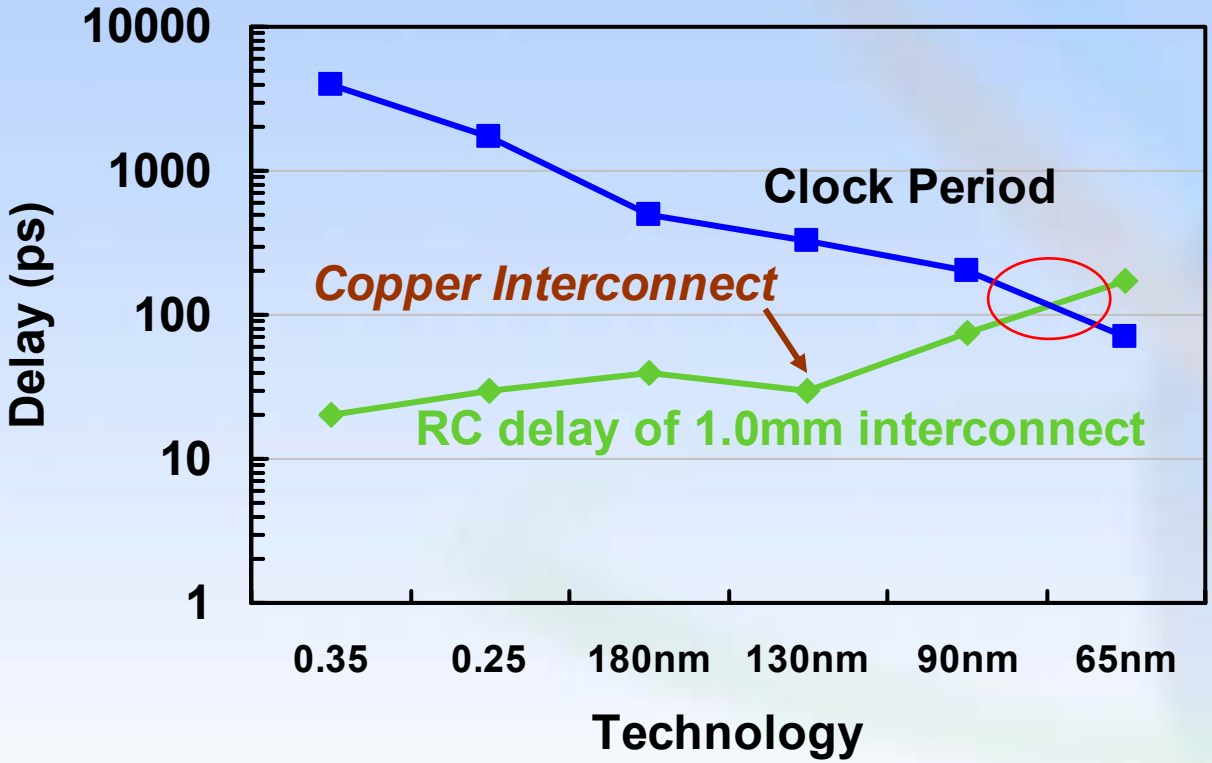

- Clock period will continue to decrease
- Delay surpassing a clock period

RC delays increasing beyond clock period

Widening Memory Gap

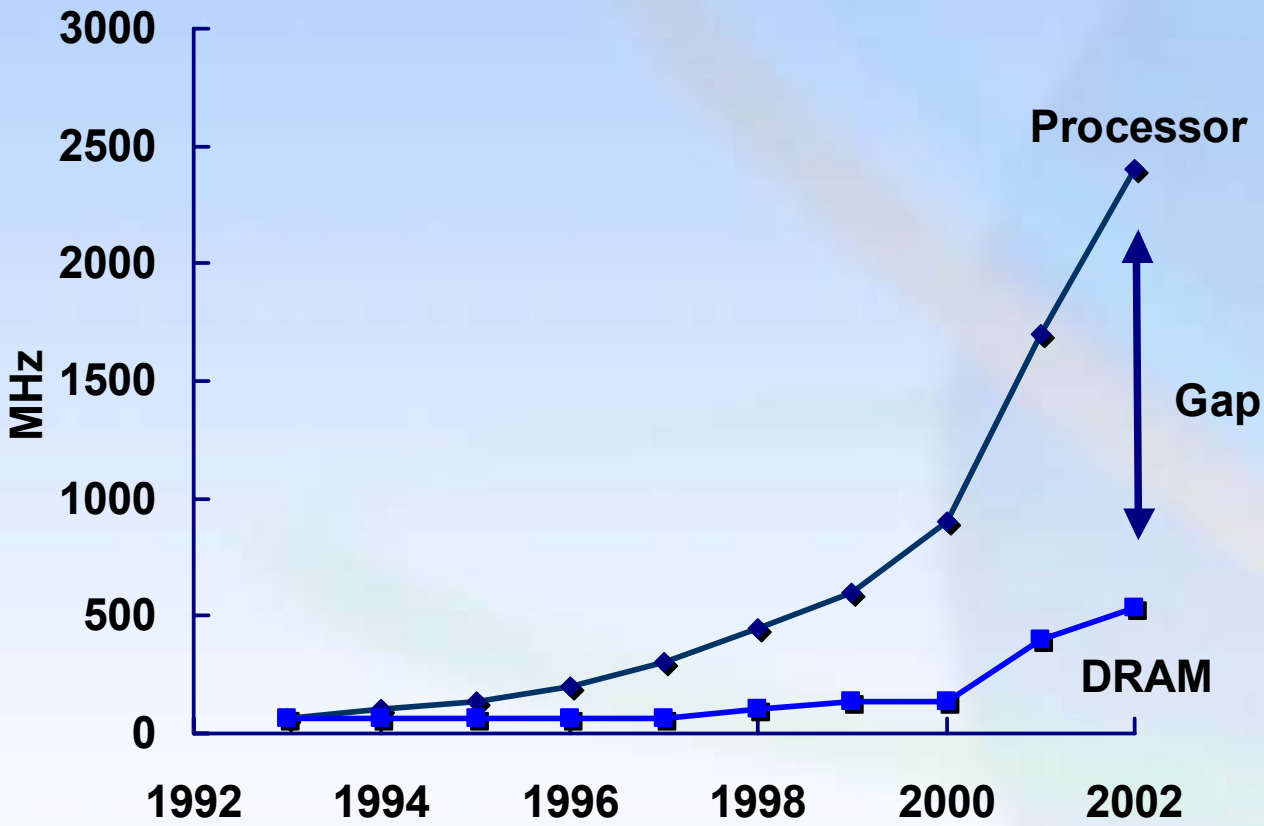

Masking the effects of slow DRAM is key

Microarchitecture Trends

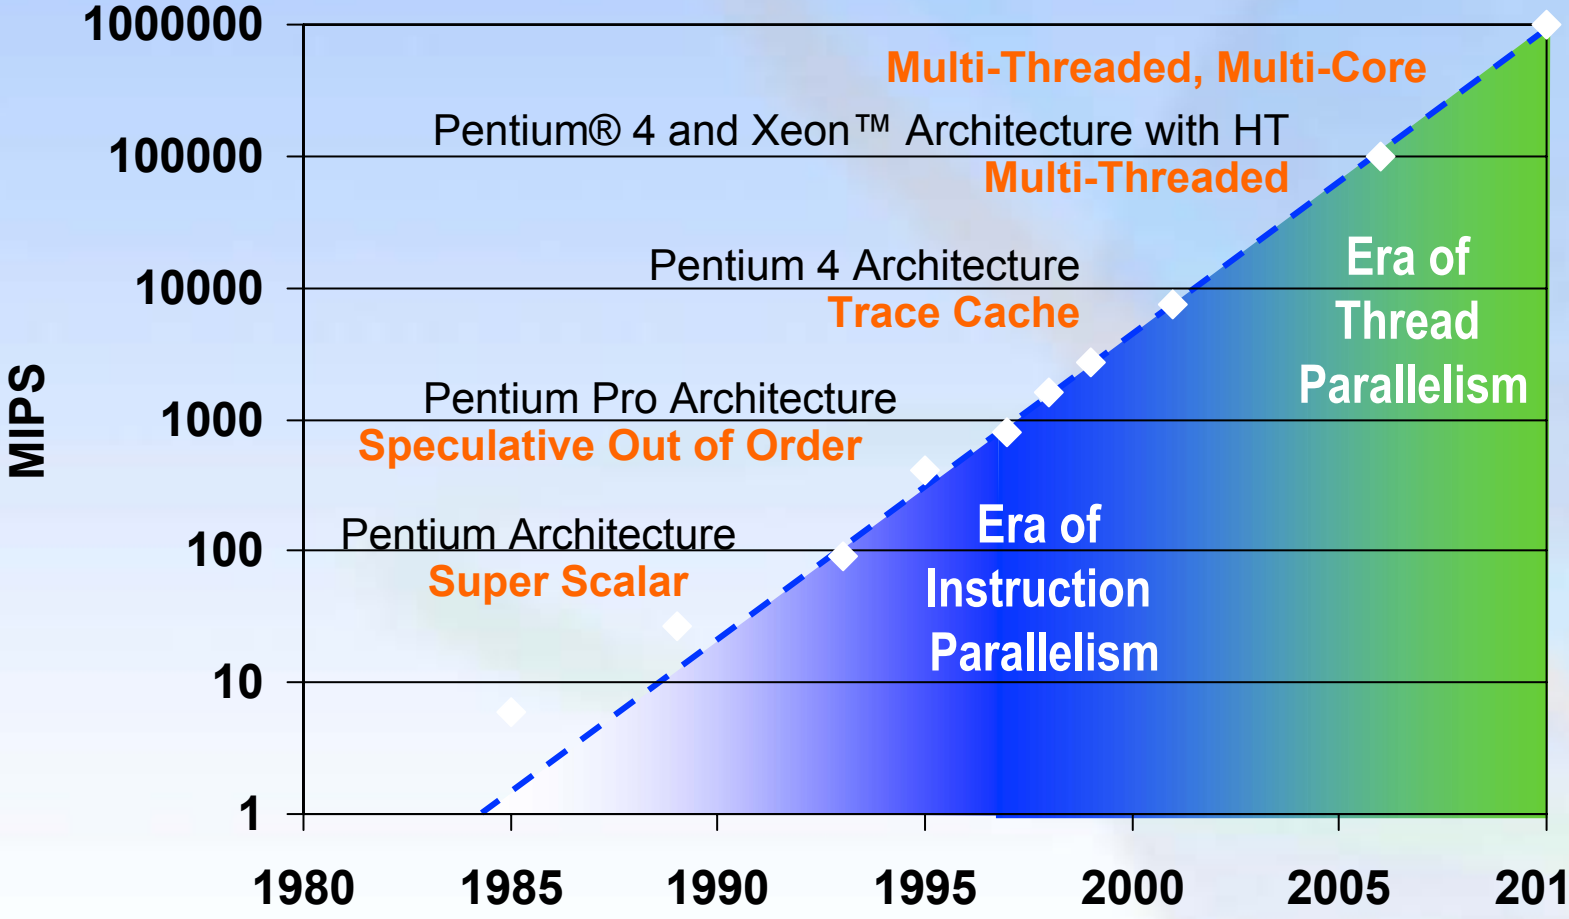

Challenge: Finding parallelism at the thread level

Agenda

- Microprocessor Trends and Challenges
- Prescott: A State-of-the-Art Microprocessor on 90nm
- Microprocessor Research
- Summary

Prescott Processor Highlights

- **Updated Microarchitecture**
  - Updated Intel® NetBurst™ micro-architecture
  - Improved Hyper-Threading Technology
  - Advanced Power Management Support
  - Prescott New Instructions
- **Larger caches**
- **Prescott is scalable to the 4-5 GHz range**
- **New Process**
  - Sub-100nm strained silicon technology
  - Seven layers of low-K metal interconnect

Updated Intel® NetBurst™ Micro-Architecture

- Improvements for higher frequency and raw performance to enable next generation applications

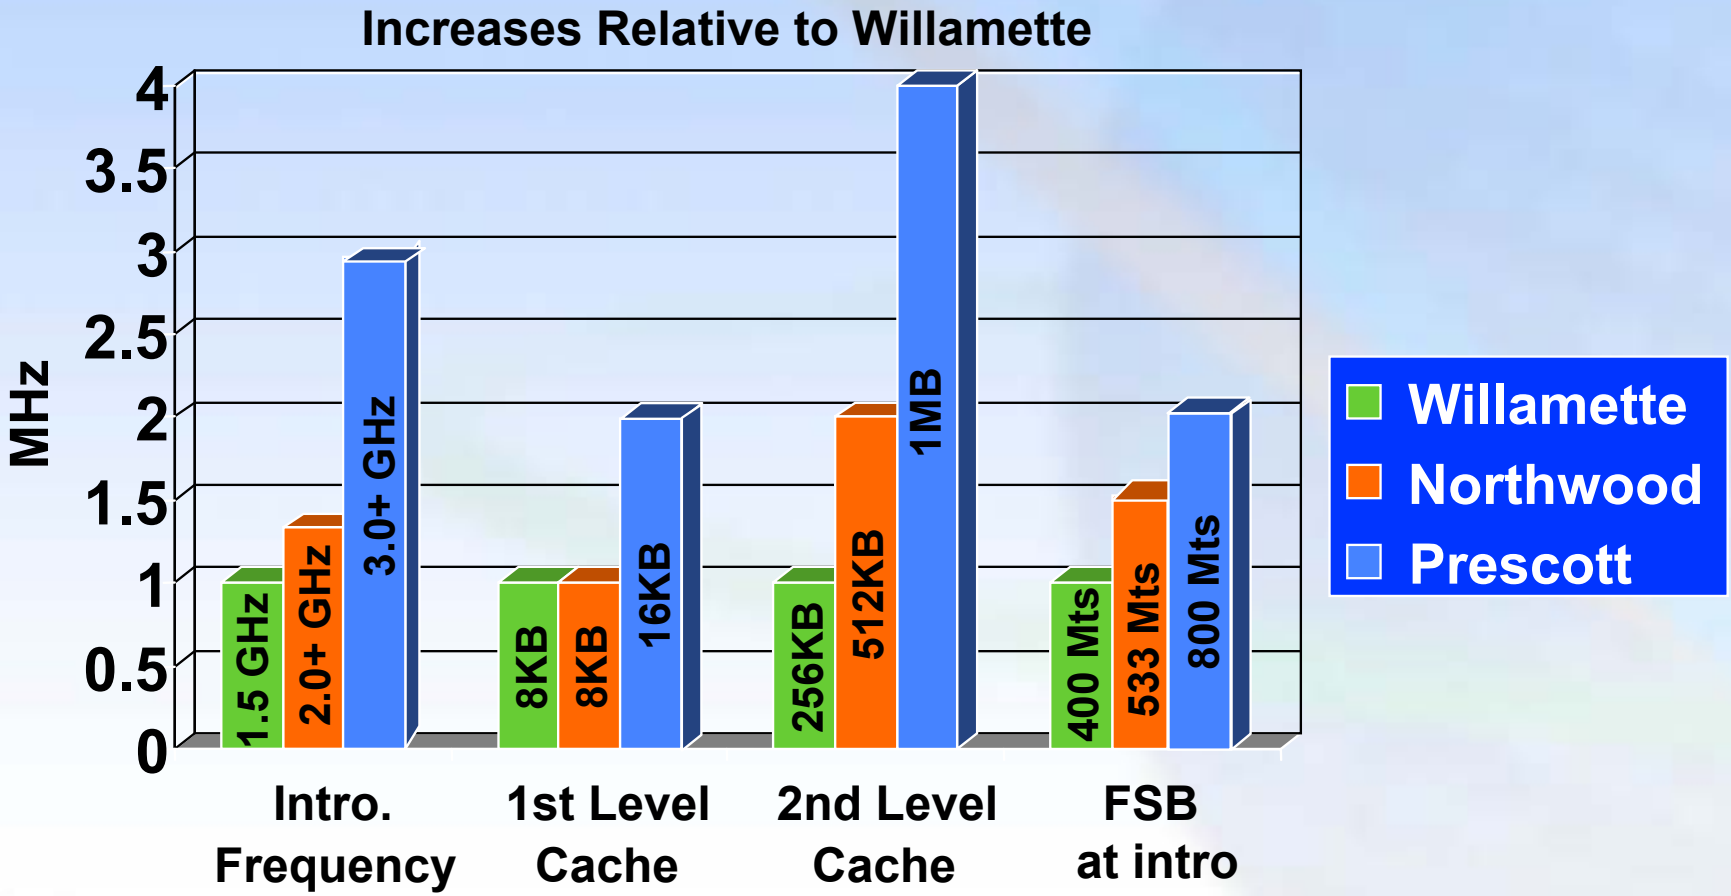

Features are for planning purposes only, and subject to change without notice.
Each 1M transfer has 8 bytes

Micro-Architecture Improvements

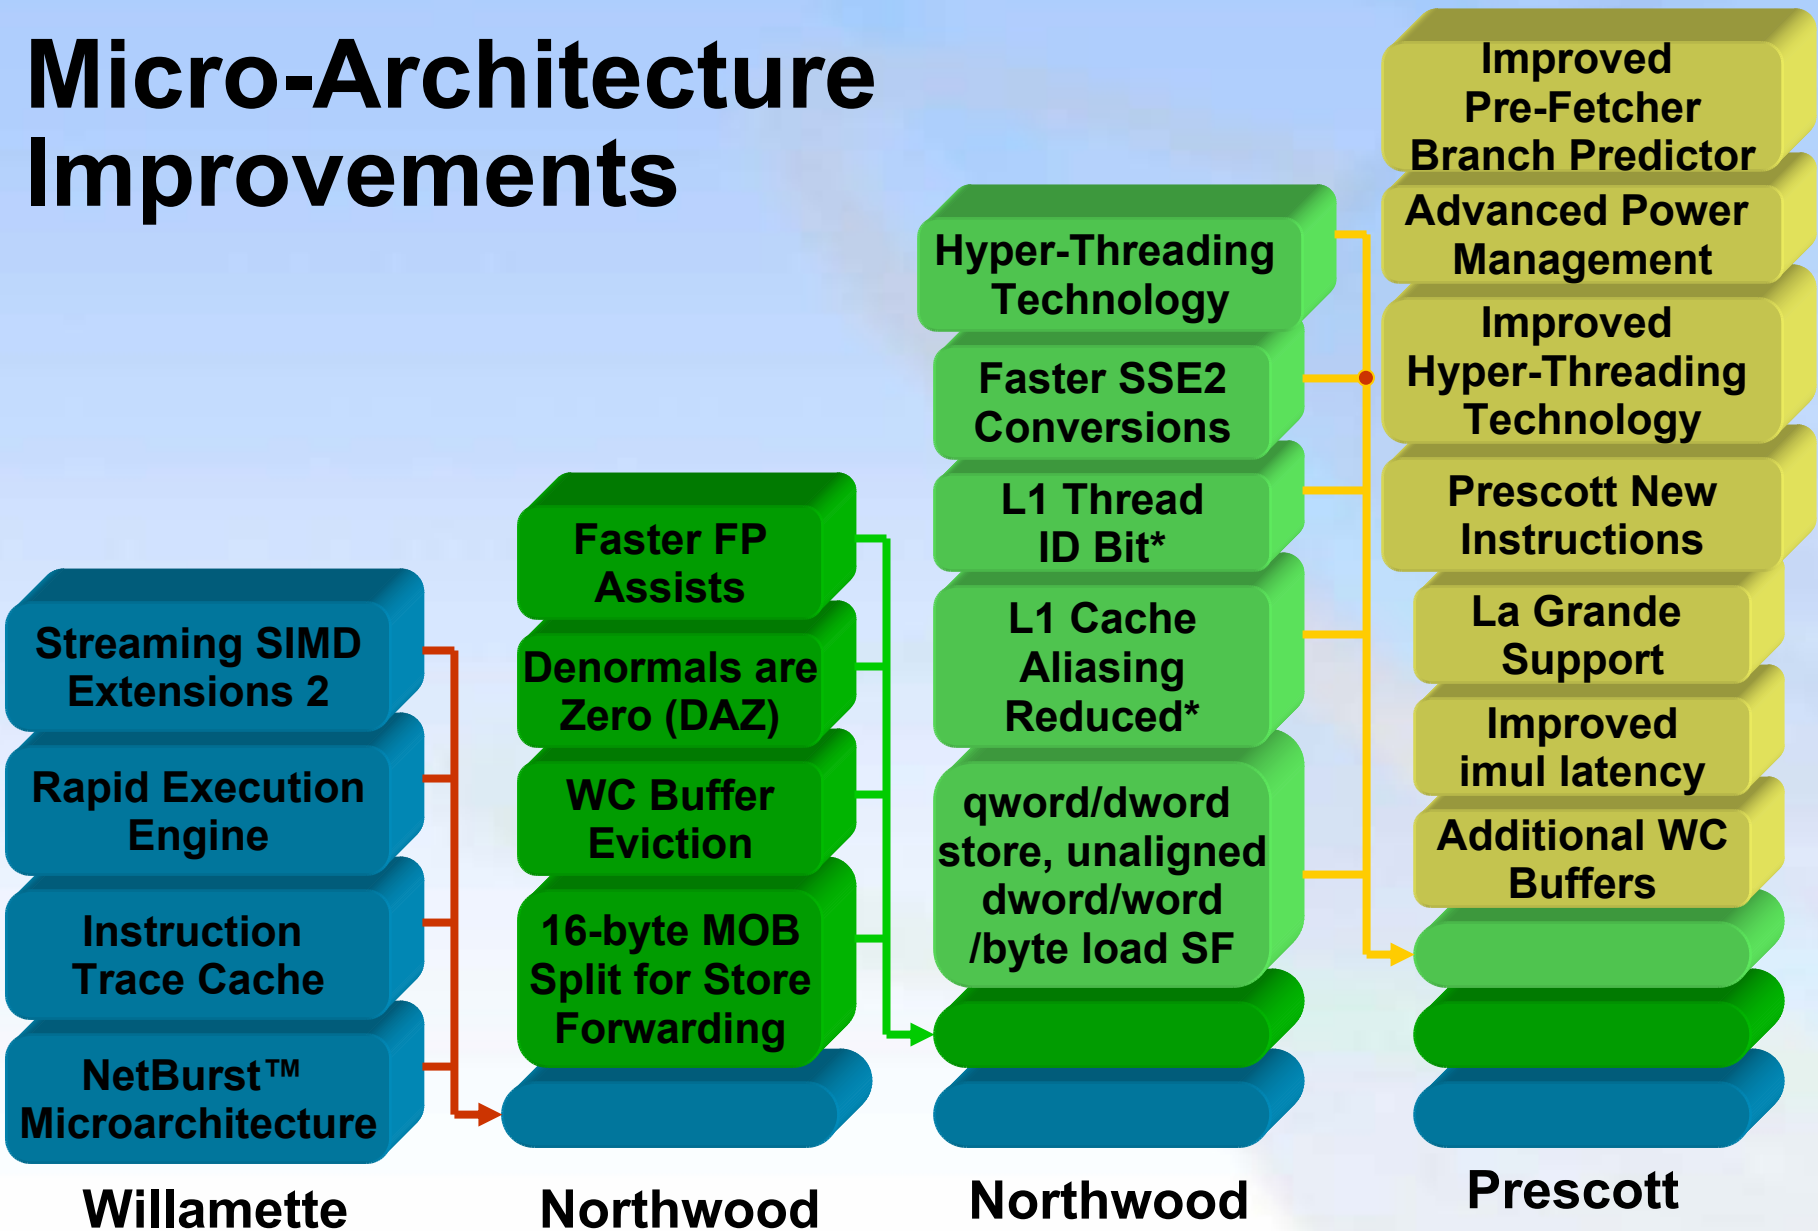

*u-arch updates which improves Hyper-Threading Technology

Prescott New Instructions

- 13 instructions added to improve specific application areas such as Media and Gaming

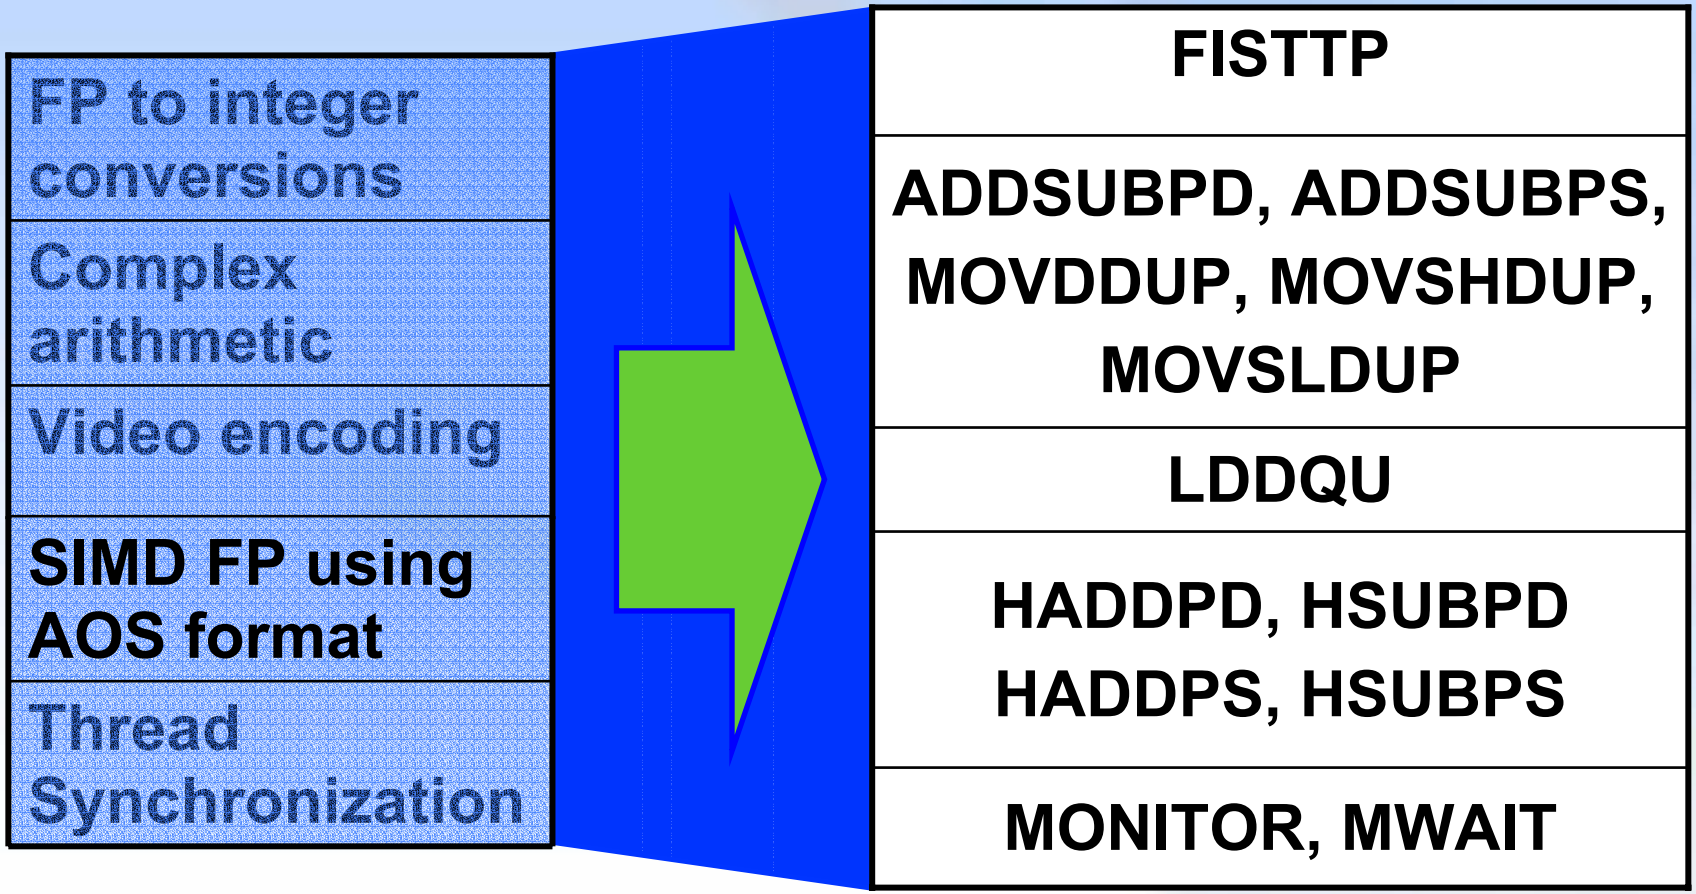

Improved Clock Distribution For Better Frequency Scaling

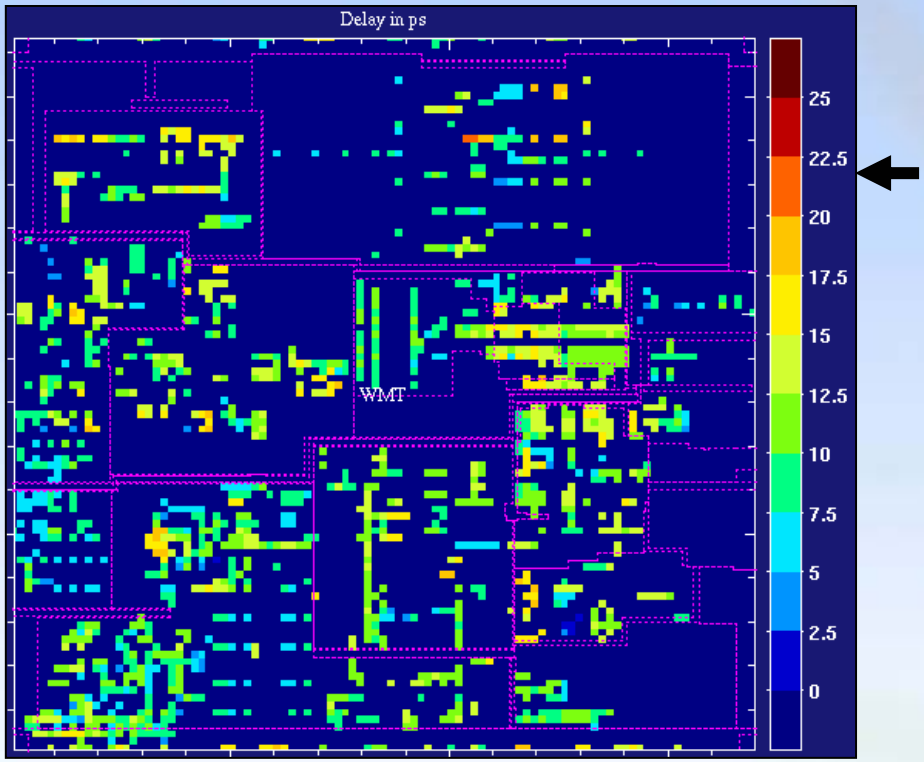

Northwood

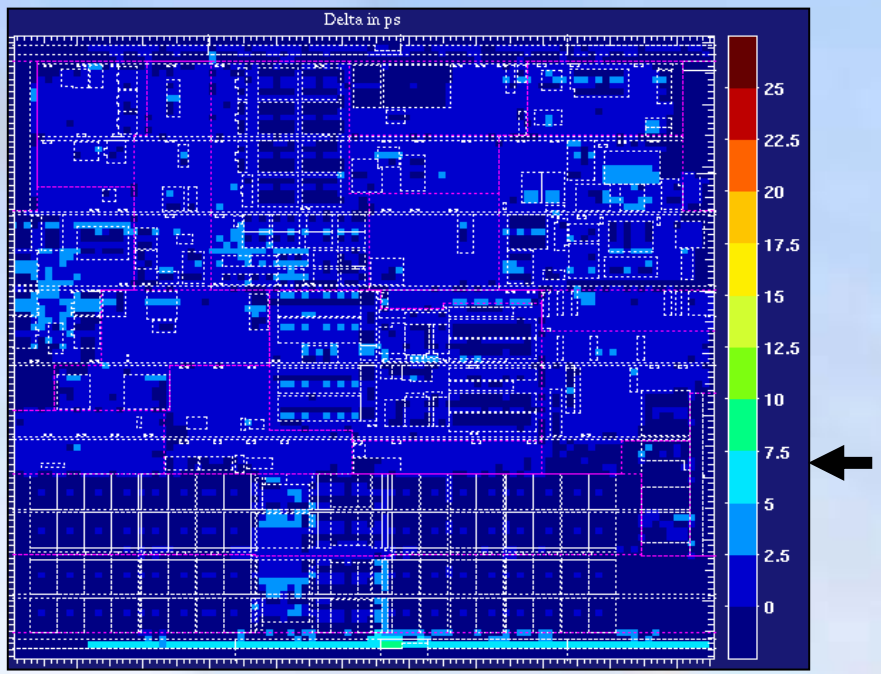

Prescott

- **Industry Leading Clocking Results**
  - 4X better than Northwood
  - Skew less than one inverter
  - Power comparable to H-Tree

90nm Performance Advantages

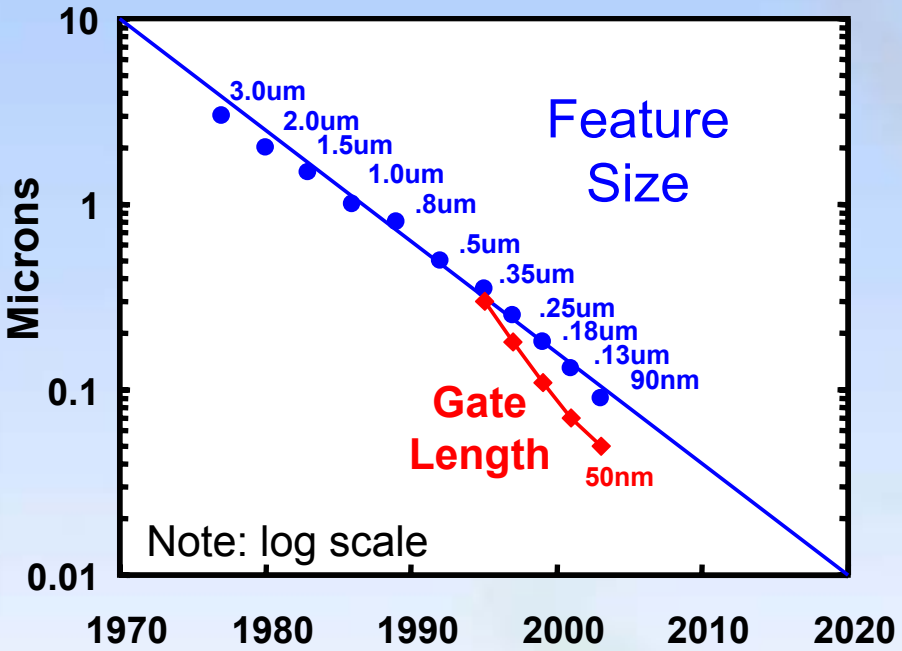

Faster gate length scaling to maintain transistor performance lead

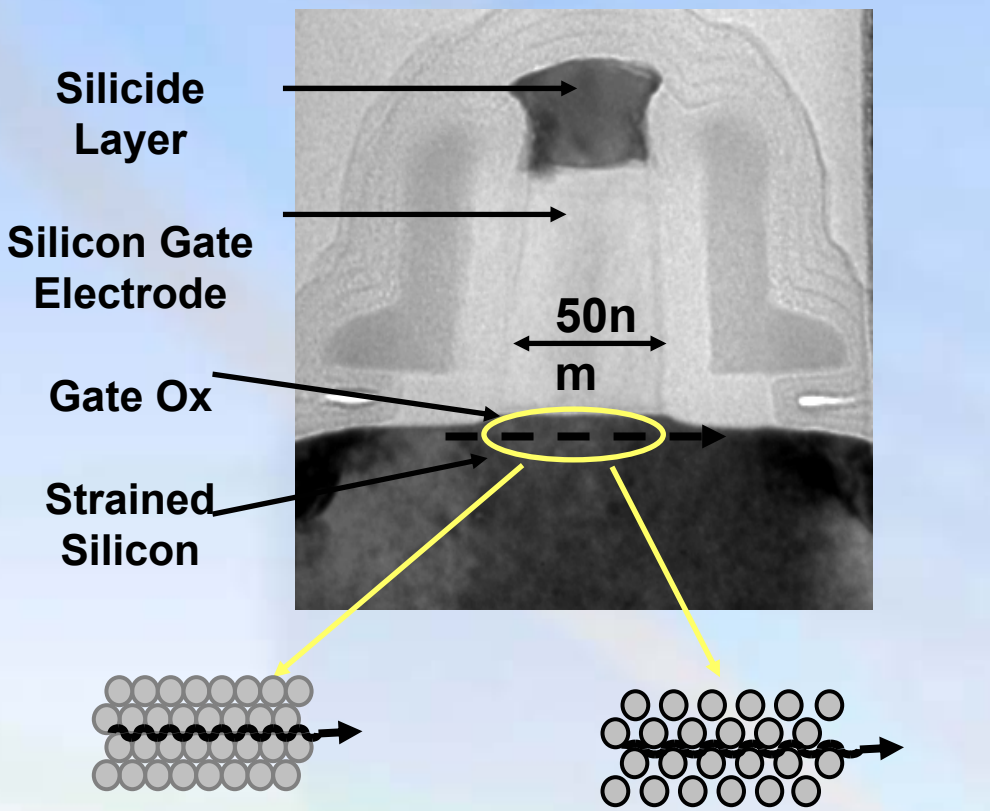

Normal Silicon Lattice

Normal electron flow

Strained Silicon Lattice

Faster electron flow

90 nm Interconnect Advantages

Low-k CDO Dielectric

New low-k dielectric introduced to reduce wire-wire capacitance

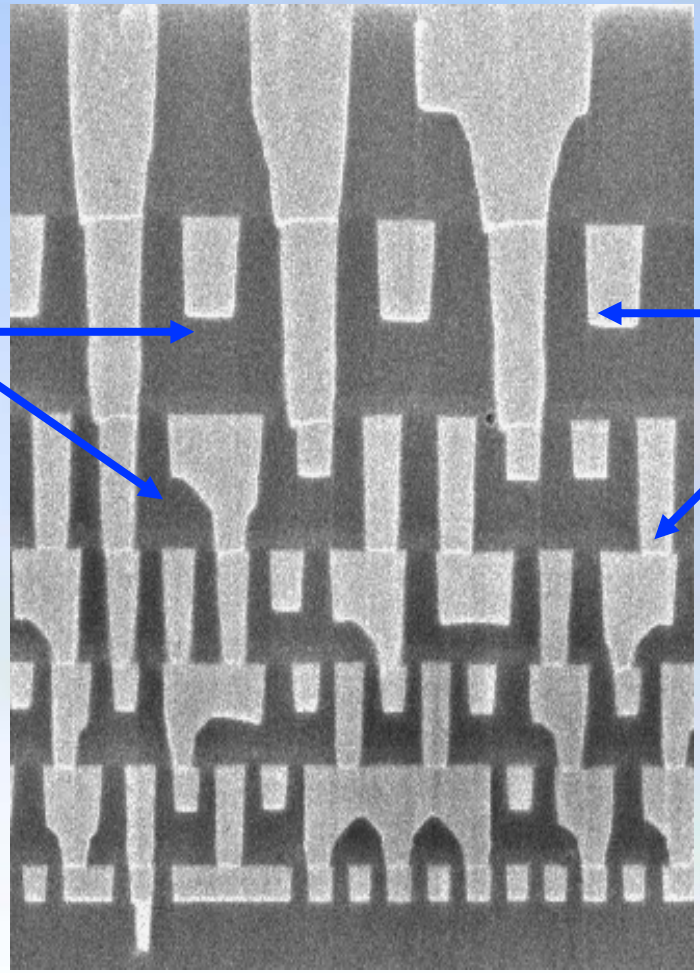

Copper Interconnects

7 layers of copper interconnect

Automation Boosts Performance and Efficiency

Northwood Floating Point

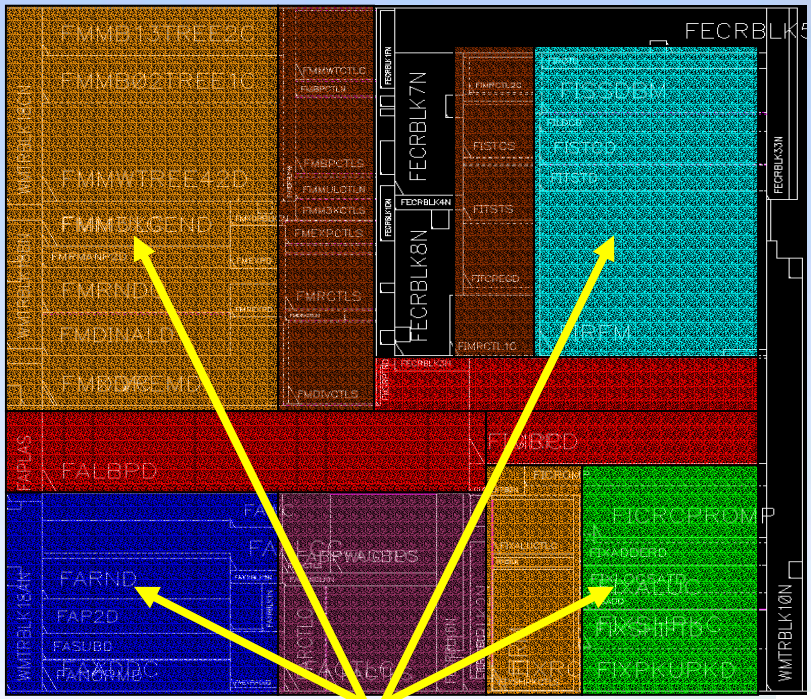

Dataflow Direction

Prescott Floating Point

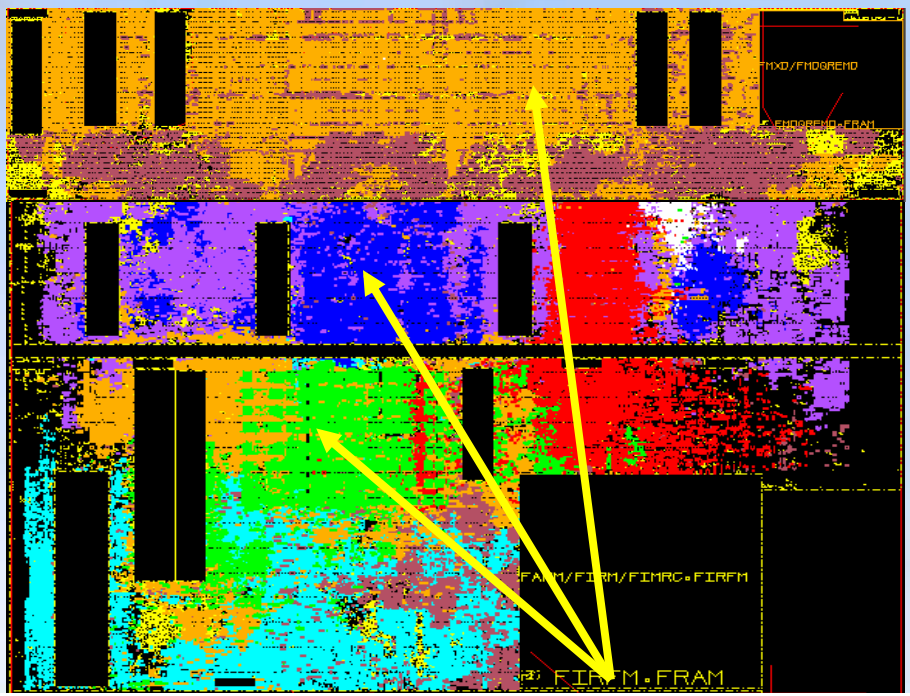

Dataflow Direction

- Superior Design Methods and Tools
  - Design completion through intelligent automation
  - Ease in building the top level chip
  - Faster design improvement with less resources

Prescott Summary

- New architecture features
  - To keep the product line on the leading edge
- New circuit techniques
  - To break scaling barriers
- New 90nm process
  - To lead in performance

Agenda

- Microprocessor Trends and Challenges
- Prescott: A State-of-the-Art Microprocessor on 90nm
- Microprocessor Research Directions
- Summary

Variations in Process, Voltage, and Temperature

Process

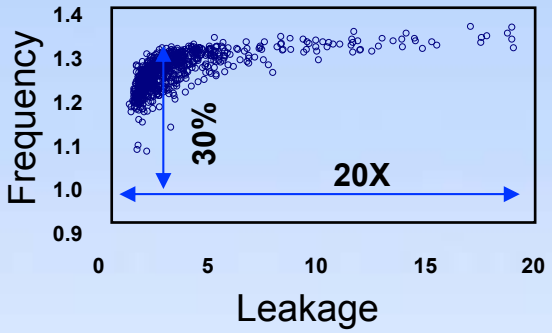

- Die-to-die variation
- Within-die variation
- Static for each die

Voltage

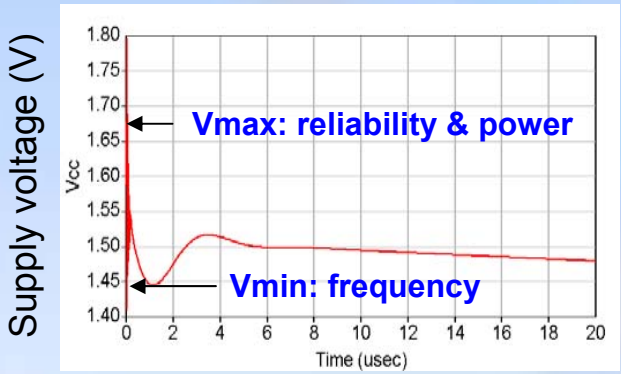

Time (usec)

- Chip activity change
- Dynamic: ns to10-100us
- Within-die variation

Temperature

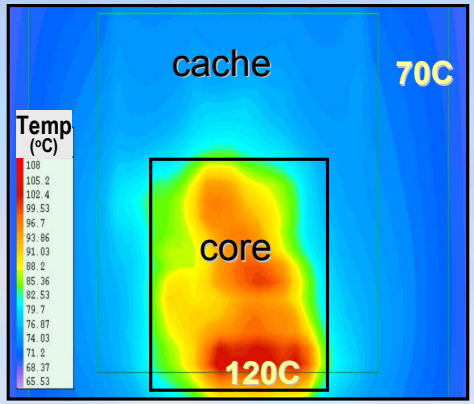

- Activity & ambient change
- Dynamic: 100-1000us
- Within-die variation

Process, voltage and temperature contribute to power and perf. variations

Sub-threshold Leakage (I_{off}) & Power

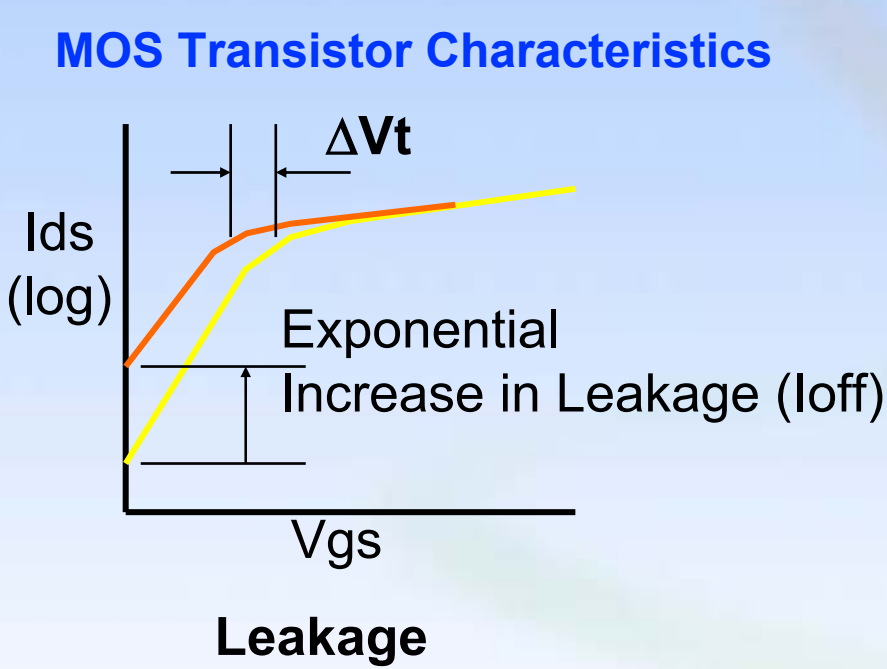

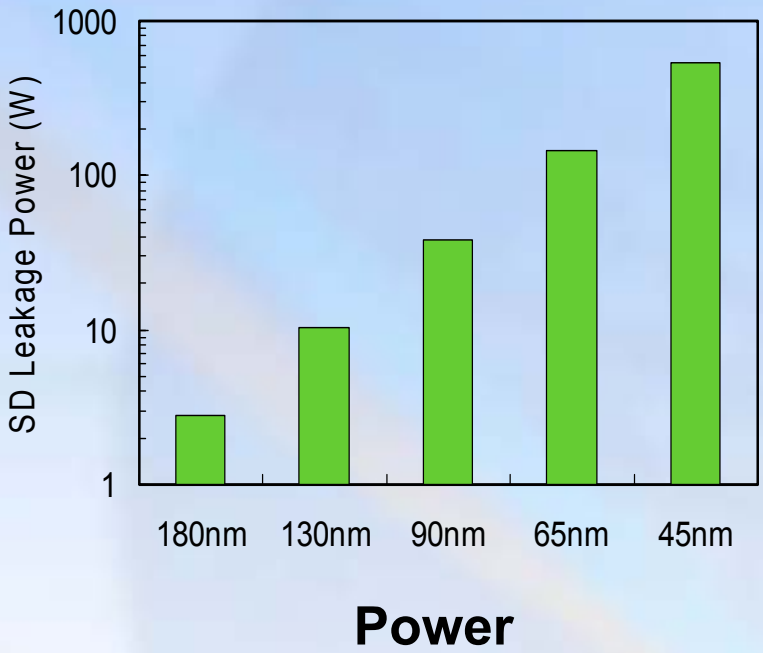

Sub-threshold Leakage and Power Increase

Two Types of Multithreading

Symmetric Hyper-Threading

Asymmetric Hyper-Threading

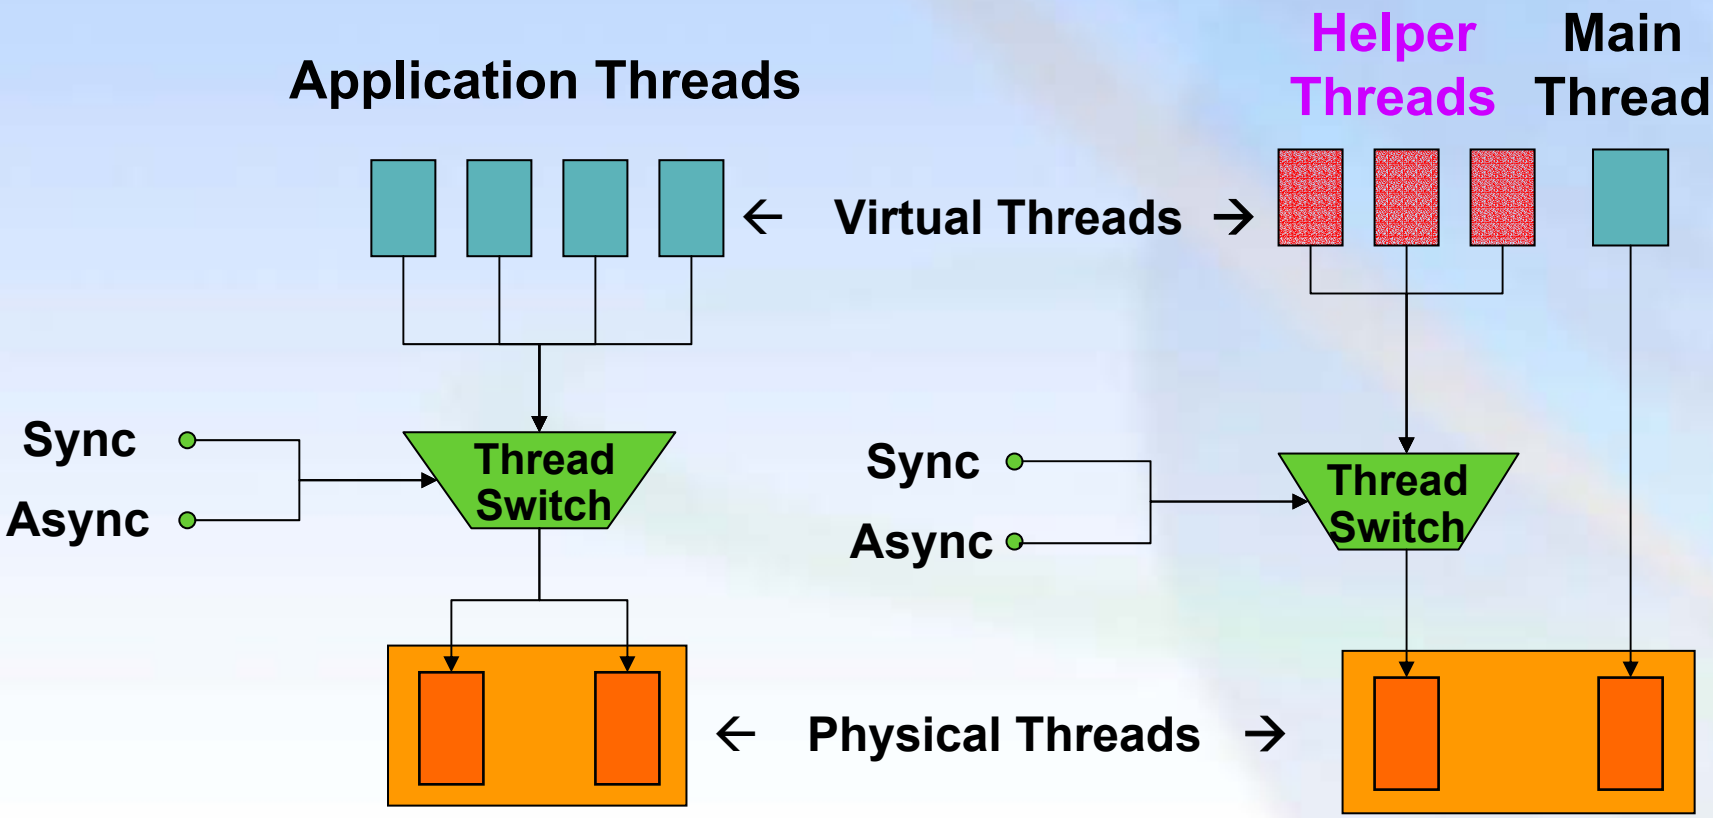

“Speculative Pre-computation”

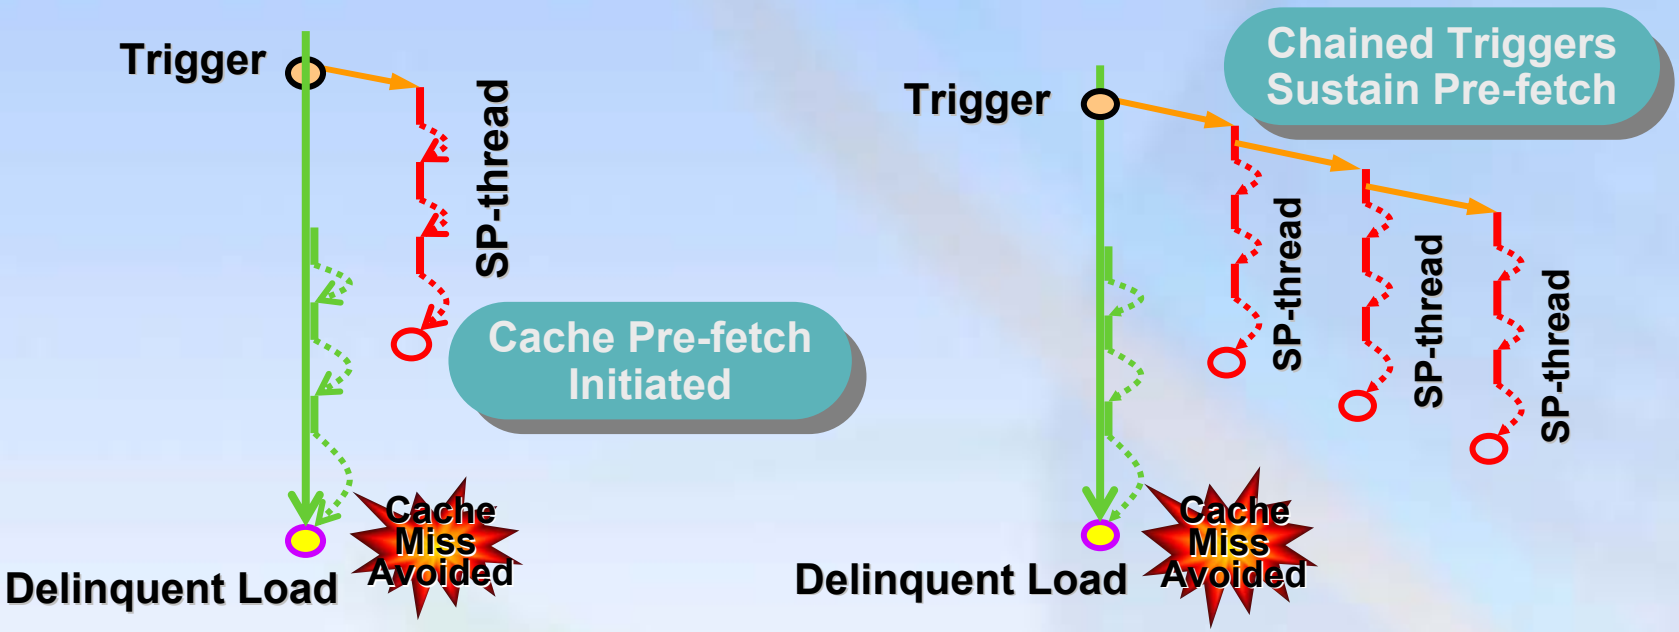

Target: The Memory Gap

Pointer-intensive applications

Pre-fetch for “Delinquent loads”

Method: SP via Helper Threads

Embed pre-fetching SP threads in binary

Parallel execution of main and SP threads

Auto-threading Compilers

- **Symmetric Hyper-Threading Compiler**
  - Partition single thread into multiple threads
  - Must ensure semantic correctness
  - Difficult for common and legacy code

- **Asymmetric Hyper-Threading Compiler**
  - Attach helper threads to original code
  - Leverage side effect of helper threads
  - Can be dynamically invoked/controlled

intel

Probabilistic Computing

- **Bayesian Networks** provide a unifying methodology in order to **reason in the presence of uncertainty**
- At the heart of many areas of computing including search, machine learning, robotics, vision…

Example: turbocodes = practical methods of communicating reliably at rates near channel capacity

or …

performing inference in a Bayesian Network

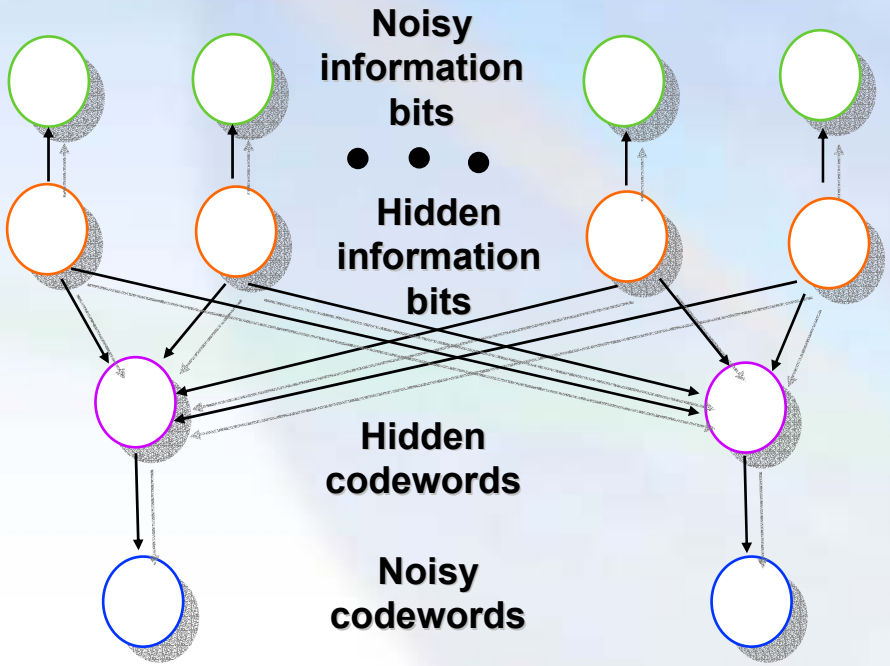

Voice & Vision Enabled Platform

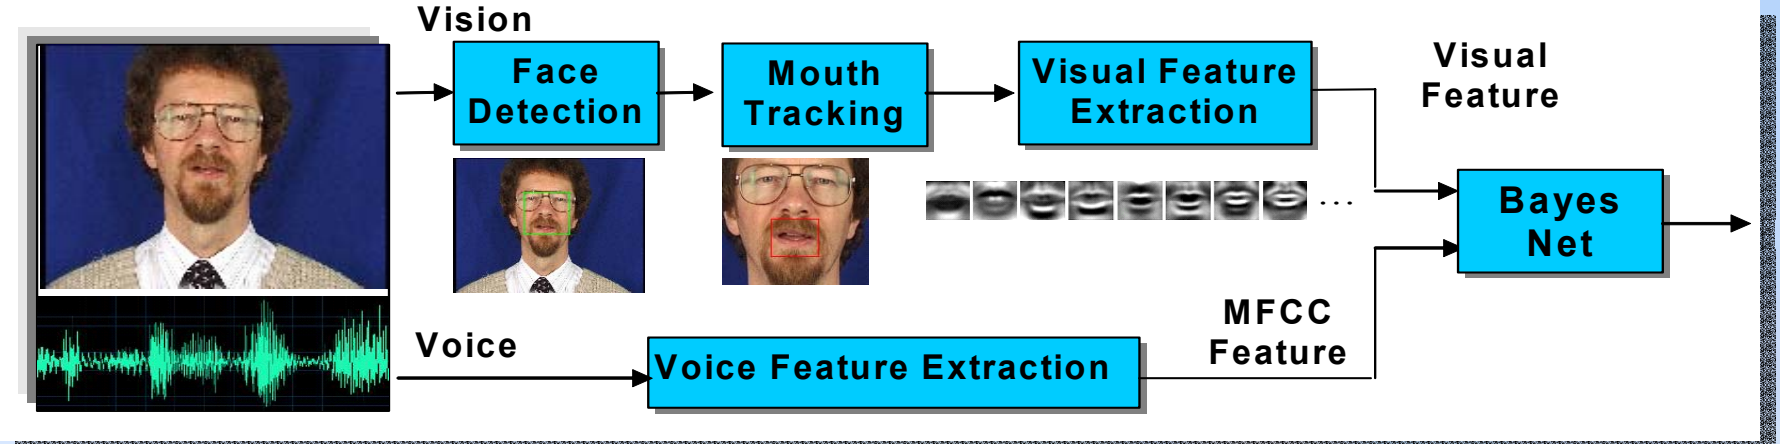

- **A/V Speech Recognition System**
  - Vision-assisted speech recognition
  - Robust in noisy environments
  - Probabilistic computing based modeling
- **Host-based voice processing**
  - Enhanced audio quality for telephony & speech
  - Beam forming, AEC, and multi-speaker tracking
  - Integrated microphone array

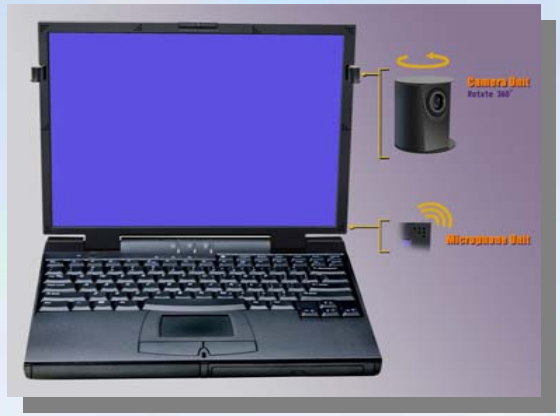

More robust 'natural' Interface

Summary

Performance Growth Continues

- Enhanced NetBurst™ microarchitecture
- Reaching 15-20 GHz by 2010

Challenge is Continued Scaling

- Difficult voltage, frequency, and leakage issues
- New circuits and new design tools

New Frontiers to Explore

- Low-leakage circuits and Hyper-Thread uarch
- New applications paradigms

Lots of ideas to explore for the next generation of processors

Please remember to fill out the session survey!